

Sacrificing information for the greater good: how to select photometric bands for optimal accuracy

Kristoffer Stensbo-Smith¹, Fabian Gieseke¹, Christian Igel^{1,3}, Andrew Zirm²,
and Kim Steenstrup Pedersen^{1,3}

¹*Department of Computer Science, University of Copenhagen*

²*Dark Cosmology Centre, Niels Bohr Institute, University of Copenhagen*

³*Space Science Center, University of Copenhagen*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

Large-scale surveys make huge amounts of photometric data available. Because of the sheer amount of objects, spectral data cannot be obtained for all of them. Therefore it is important to devise techniques for reliably estimating physical properties of objects from photometric information alone. These estimates are needed to automatically identify interesting objects worth a follow-up investigation as well as to produce the required data for a statistical analysis of the space covered by a survey. We argue that machine learning techniques are suitable to compute these estimates accurately and efficiently. This study promotes a feature selection algorithm, which selects the most informative magnitudes and colours for a given task of estimating physical quantities from photometric data alone. Using k nearest neighbours regression, a well-known non-parametric machine learning method, we show that using the found features significantly increases the accuracy of the estimations compared to using standard features and standard methods. We illustrate the usefulness of the approach by estimating specific star formation rates (sSFRs) and redshifts (photo- z 's) using only the broadband photometry from the Sloan Digital Sky Survey (SDSS). For estimating sSFRs, we demonstrate that our method produces better estimates than traditional spectral energy distribution (SED) fitting. For estimating photo- z 's, we show that our method produces more accurate photo- z 's than the method employed by SDSS. The study highlights the general importance of performing proper model selection to improve the results of machine learning systems and how feature selection can provide insights into the predictive relevance of particular input features.

Key words: galaxies: distances and redshifts – galaxies: star formation – galaxies: statistics – methods: data analysis – methods: statistical – techniques: photometric

1 INTRODUCTION

High-resolution spectroscopic data contain a wealth of information about astrophysical objects. Analyses relying on spectroscopy suffer, however, from small sample sizes. Photometric surveys have the potential to overcome this limitation, but are limited in terms of the amount of information that can be extracted for each astrophysical object. Due to the abundance of data currently available, and especially with the surveys commencing within the next decade, methods are required that can automatically extract relevant information from the broad-band images of these surveys. Our goal is to reliably, efficiently and accurately estimate properties of objects from photometric data, for example, for quickly identifying interesting objects worth a follow-up investigation or for conducting large-scale statistical analyses.

In this study, we apply a method for selecting the most informative colours and bands for photometric estimations. We illustrate its potential by estimating specific star-formation rates (sSFRs) and photometric redshifts (photo- z 's) from available SDSS data, but the method can readily be applied to other quantities and surveys.

1.1 Star formation rates

An ongoing quest in cosmology is the understanding of galaxy formation and evolution. A crucial part here is to understand the star formation history of the individual galaxies as well as the universe as a whole. Major open questions include which processes trigger star formation and, equally important, quench it. Data from large surveys, such

as the Sloan Digital Sky Survey (SDSS, York et al. 2000), have shown a peculiar bimodality in the star-formation rates (SFRs) of galaxies (Kauffmann et al. 2003). The bimodality points to a scenario where star formation is quenched, but the responsible mechanism is far from understood. Current results indicate that the quenching time-scale varies significantly with galaxy mass (Wetzel et al. 2012, 2013; Wheeler et al. 2014) and redshift (Balogh et al. 2016), suggesting that different processes are in play at different times and masses (Fillingham et al. 2015; Wetzel et al. 2015). To uncover these processes, it is natural to turn to the statistical properties of a large number of galaxies in order to look for correlations between SFRs and other physical properties.

The most common way to estimate the recent SFR of a galaxy is to use a number of observational tracers. These tracers often rely on observations of single or multiple emission lines, with the H α emission line being among the most popular (Kennicutt & Evans 2012). A main limitation is that they usually require high-quality spectra. Other methods to estimate the SFR include conversion factors to convert from flux over a given wavelength interval (Kennicutt 1998; Kennicutt & Evans 2012) and spectral energy distribution (SED) fitting, which relies on a library of template spectra generated by stellar population synthesis models (for recent reviews, see Walcher et al. 2011; Conroy 2013). In the most basic version of this method an observed galaxy spectrum is compared to every template spectrum, the closest match is chosen and the template’s physical properties adopted (e.g., Charlot et al. 2002; Brinchmann et al. 2004).

SED fitting is often considered a less precise way to estimate SFRs than relying on observational tracers (Walcher et al. 2011). It does, however, allow us to estimate the SFR from broad-band photometry, where observational tracers have more limited use (Maraston et al. 2010).

More direct estimations of SFRs and specific SFRs (sSFRs) from broad-band photometry have also been investigated (e.g., Williams et al. 2009; Arnouts et al. 2013), though there are still significant discrepancies between these estimated quantities and those obtained from more reliable methods.

1.2 Photometric redshifts

Spectroscopic surveys provide highly accurate redshifts of galaxies, enabling a detailed 3D view of galaxy distribution in the universe, but they are both expensive and time-consuming. Photometric surveys, on the other hand, can cover a much larger area of the sky in less time, and can usually go below the spectroscopic flux limit. They therefore provide a significantly more complete, and thus less biased, sample of galaxies, which is a notable advantage over spectroscopic surveys. Photometric surveys, however, struggle with reduced accuracy in the galaxy positions along the line of sight. Despite this problem, the larger galaxy sample sizes are useful for numerous cosmological applications, such as obtaining constraints on cosmological parameters (e.g., Padmanabhan et al. 2007; Carnero et al. 2012; Ho et al. 2012). These applications rely on photometric redshifts (photo- z ’s) calculated from broad-band photometry. Naturally, increasing the accuracy of photo- z ’s is of great importance.

A vast amount of methods have been developed to es-

timate photo- z ’s (see, e.g., Hildebrandt et al. 2010; Abdalla et al. 2011, for recent comparisons). Broadly speaking, photo- z estimation methods can be classified as either template-based or empirical methods. Template-based methods use SED fitting in the same way as for SFR estimation: they match the observed colours or magnitudes to those of a large library of synthetic template spectra (e.g., Benítez 2000; Bolzonella et al. 2000; Ilbert et al. 2006; Brammer et al. 2008). Empirical methods train algorithms to estimate photo- z ’s from colours or magnitudes. The algorithms are calibrated to fit the task at hand using a training dataset with spectroscopically derived redshifts.

A wide range of empirical methods have been developed, and most fall into the categories of either tuning the colour- z relation or machine learning. The machine learning category is highly diverse, with techniques such as artificial neural networks (Collister & Lahav 2004), self-organising maps (Geach 2012), random forests (Carrasco Kind & Brunner 2013), and Gaussian processes (Almosallam et al. 2016) having been used for photo- z estimation. These techniques generally outperform template-based methods for photo- z estimation, as machine learning methods are able to adapt to the highly nonlinear relation between colours and redshift. For recent reviews of the performances of various photo- z estimation methods, see Dahlen et al. (2013) and Sánchez et al. (2014).

1.3 Increasing the information from photometric measurements

SED fitting is a common method for both photo- z and SFR estimation. Advantages of this method include the ability to get the full star formation history (SFH) (limited by the detail level of the template library) of a galaxy as well as constraints on its redshift, environment etc. The restrictions lie in the generation of the template spectra, with computational power and understanding of stellar evolution being the main limiting factors.

The main computational limitation is the enormous amount of free parameters that can be tweaked in the generation of a single spectrum. Because of this, and limited physical knowledge about stellar evolution, it is still a great challenge to generate appropriate template spectra (e.g., Pacifici et al. 2015; Smith & Hayward 2015). A brute-force way of calculating templates for a chosen grid of parameters quickly becomes infeasible. The amount of degeneracies between the evolutionary states of different single stellar populations (SSPs) also limits this approach.

A number of ways to reduce the amount of necessary template spectra with minimum information loss have been explored. In particular, machine learning methods have been used to interpolate between template spectra to allow for a sparser grid to be sampled (e.g., Tsalmantza et al. 2007). Active learning was explored by Solorio et al. (2005), where the computer automatically generates new template spectra if no close match is found in the dataset. This automatically refines the template grid in regions that have actual observations. A different approach was taken by Richards et al. (2009), who used diffusion K -means to tackle the problem of choosing which SSPs make up a galaxy spectrum, by finding an appropriate basis from a large set of SSP spectra. In the same spirit, Chen et al. (2009) used a principal component

analysis to estimate specific sSFRs from obtained eigenspectra.

While spectroscopy is superior in terms of information content, photometry excels in terms of coverage. Using a machine learning approach to estimate parameters can give us the best of both worlds. The algorithm can be trained on galaxies with accurate parameters determined from high-resolution spectra and then be used to estimate the same parameters of other galaxies from broad-band photometry only. This avoids the problem of generating template spectra from models that may suffer from various restrictions and approximations. However, just as template-based methods require the parameter space to be densely sampled in order to provide good parameter estimations, machine learning methods require training data that represent the entire population. If such are not available, the methods may lead to biased estimates. Machine learning methods can also achieve significantly lower computational complexity compared to SED fitting, depending on the level of detail wanted, which will become increasingly important in the near future, when new photometric surveys start producing data at an unprecedented rate.

Using highly detailed data can, however, lead to a decrease in accuracy. This counter-intuitive phenomenon occurs for both template methods as well as machine learning methods, and can be attributed to the fact that if a dimension contributes only (or even just some) noise, it will decrease the overall signal-to-noise ratio (S/N).

Selecting only the most informative dimensions of the data can therefore lead to higher accuracy, even if it requires removing somewhat informative dimensions, as the lower dimensionality of the data can result in a higher S/N.

In the machine learning literature, the dimensions of a data point are referred to as *features*. Thus, the task of choosing the most informative dimensions is called *feature selection*. Feature selection has already been investigated in an astrophysical context. Among the most used feature selection algorithms are random forests, which produce feature ranking as part of the algorithm. They have been used in a number of studies, for example, D’Isanto et al. (2016) and Rimoldini et al. (2012). Random forests are not the only way to select features, and Graham et al. (2013) tested five different feature selection strategies for classifying stars. Hoyle et al. (2015) showed how adding the most informative features to the standard set of colours and magnitudes significantly increased the accuracy for photo- z estimation.

It is important to realise that the concept of most informative features is not a universal one; the most informative features for one algorithm may be different from those of another. That depends on how specifically the algorithm uses the features, for example, some algorithms may be sensitive to scaling of the features, while others may not. And just as the most informative features vary from algorithm to algorithm, so will they vary from task to task. For example, whereas observed UV radiation may contain a lot of information regarding star formation in the nearby universe, it may not be that informative for detecting, say, brown dwarfs.

In this paper, we show that we can obtain a significantly greater accuracy of estimated photo- z ’s and sSFRs, using only SDSS *ugriz* photometry, by applying a machine learning method rather than relying on spectral modelling of the photometry. Our approach is similar to that of

Smidt et al. (2013), but here we show that the accuracy can be further increased by performing a feature selection, selecting the most informative features among all measured SDSS magnitudes and colours.

Specifically, we use k -nearest neighbours (k -NN) regression, which is an intuitive method well-known in machine learning and to some extent also in astronomical communities (see, e.g., Li et al. 2008; Polsterer et al. 2013, 2014; Kügler et al. 2015; Kremer et al. 2015). Of the more prominent uses of k -NN in astronomy is the estimation of photo- z ’s in SDSS (Abazajian et al. 2009).

By using k -NN we can automatically learn a mapping from magnitudes and colours of galaxies to their parameters derived from reliable indicators, thereby allowing accurate photometric estimates without high-resolution spectra. The reliable parameters can be estimated using any method deemed appropriate for each individual galaxy, effectively taking advantage of multiple indicators, as explored by Wuyts et al. (2011, 2013) for SFRs. A significant advantage of k -NN over other methods is that it naturally adapts to the local, potentially high-dimensional structure of the data, and can thus model highly non-linear behaviour without problems. Another virtue of k -NN is its simplicity, which makes it easy to see how data are used and compared within the algorithm.

Selecting the most informative features can, in theory, be done by trying all possible feature combinations. As the number of combinations grows exponentially with the number of features, this quickly becomes unfeasible, and one has to resort to clever selection strategies. Here, we use *forward feature selection* to determine the most informative features (see section 2.3 for details). Forward feature selection was used by Xu et al. (2013) to examine which halo properties contained most information about the number of galaxies. In this paper, we use it to improve the estimation of photo- z ’s and photometric sSFRs, which illustrate the method’s general usefulness.

The remainder of this paper is organised as follows: in section 2 we describe the k -NN algorithm and the algorithm we use to select the most informative colours. Section 3 describes the data we are using and details our experimental set-up. In section 4 we provide results of our experiments and an analysis of these. We end with a discussion and a summary of our conclusions in section 5.

2 METHODS

The goal of this study is to test the efficiency of machine learning techniques, in particular feature selection, when estimating physical quantities of galaxies. We suggest using the selected features directly in regression methods rather than in connection with physical models, such as population synthesis models. There are two fundamental ways of doing regression: parametric and non-parametric. In the parametric case, data is assumed to follow a function $f(x)$ with known form but unknown parameters. It is usually fairly easy to estimate these parameters by fitting, but this advantage comes at a cost: by choosing a particular functional form of $f(x)$ we have made assumptions about the underlying structure of the data. If these assumptions are not absolutely correct, we will not be able to achieve optimal

estimation performance (James et al. 2013). This is where non-parametric methods have an advantage, as they do not make any assumptions about the structure of the data, but adapt to it.

2.1 k nearest neighbours regression

We employ one of the simplest non-parametric methods, namely k nearest neighbours (k -NN) regression (Altman 1992; Hastie et al. 2009; James et al. 2013). Assume that we are given a data set $S = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N)\} \subset \mathbb{R}^D \times \mathbb{R}$ consisting of D -dimensional data points \mathbf{x}_i with associated output values y_i . For instance, each data point could represent a galaxy with $D = 2$ colour values (e.g., $B - V$ and $U - B$) and the output value y_i could be the sSFR that one is interested in estimating. The components of \mathbf{x}_i (which, in this example, would be the colours) are called *features*. Now we employ machine learning to infer from S a general rule of how to predict the (unknown) output value y' given some new data point \mathbf{x}' . The k -NN method does this by simply finding the k closest data points with known output values, and then taking the average of these values, i.e.,

$$y' = \frac{1}{k} \sum_{i \in \mathcal{N}_k} y_i, \quad (1)$$

where \mathcal{N}_k is the set of the k nearest data points in $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ w.r.t. the new sample \mathbf{x}' . The ‘closeness’ between samples is defined via a metric d . That is, $\mathcal{N}_k = \mathcal{N}_{k-1} \cup \arg\min_{(x,y) \in S \setminus \mathcal{N}_{k-1}} d(\mathbf{x}, \mathbf{x}')$ for positive integers k and $\mathcal{N}_0 = \emptyset$, where $\arg\min$ breaks ties at random. We use the Euclidean metric $d(\mathbf{x}, \mathbf{z}) = \sqrt{\sum_{i=1}^D (x_i - z_i)^2}$ for $\mathbf{x}, \mathbf{z} \in \mathbb{R}^D$, though any metric can be chosen. The Euclidean distance is the most common choice in the literature, but it is perfectly possible that another metric would perform better. One can also attempt to learn the metric from the data as done by, e.g., Weinberger & Saul (2009). To keep things simple, however, we stick to the Euclidean metric.

Although the k -NN regression method is simple, it often yields highly accurate predictors. This is especially the case if the amount of training data N is large and/or the feature space dimensionality D is low. While it may seem counterintuitive, adding more features (i.e., dimensions) to the input data may make k -NN perform worse. The performance of nearest neighbours models can deteriorate if D gets too large, in particular when each added dimension contains intrinsic noise. The addition of extra noise with each added dimension may eventually decrease the S/N. This is perhaps most easily recognised if one considers the extreme case of adding a feature, which is pure noise. This can only decrease the performance, and adding more of these pure noise features will eventually down any signal present in the original features. Thus, it is important to select the right features, see section 2.3.

2.1.1 Dealing with uncertainties

In its most basic form, the k -NN algorithm does not support the inclusion of uncertainties associated with inputs or outputs, nor does it provide confidence intervals for the estimated quantities beyond calculating the variance of the

neighbours’ outputs (Altman 1992). There are, however, extensions dealing with these issues.

There are a number of ways uncertainties may influence the results of an analysis. Firstly, there may be uncertainties related to the output values (e.g., sSFRs or photo- z ’s) of the training data, which need to be propagated to the predicted output. Secondly, there may be uncertainties in the input values (e.g., colours) of both the training data and the new data, which also need to be propagated to the estimated output value.

Propagating uncertainties from known data to the estimate done by k -NN is not a trivial task. Ideally, to estimate the output value of a new datum, its input uncertainties need to be propagated, and one needs to incorporate the uncertainties on both input and output of the training data. A standard Monte Carlo sampling can deal with all these uncertainty issues, but it will quickly get far too computationally expensive.

Assuming Gaussian errors, uncertainty in the output alone can be dealt with in a relatively straightforward manner by using a weighted average, $y' = \sum_{i \in \mathcal{N}_k} w_i y_i / \sum_{i \in \mathcal{N}_k} w_i$, using $w_i = \sigma_i^{-2}$, where σ_i^2 is the variance of y_i . This does not account for the scatter of the inputs, which ideally should mean that more distant neighbours (and their corresponding uncertainties) are weighted less when computing the average. This can be accounted for by including the similarity metric in the weights, or, alternatively, including the uncertainties in the similarity metric as done by, e.g., Polsterer et al. (2013). An additional complication arises due to the uncertainties in the inputs and the choice of number of neighbours, k . With uncertain inputs, the question of which of two neighbours is closer cannot be answered with complete certainty.

Both the question regarding choosing k and that of choosing the proper similarity metric can, however, be addressed with a probabilistic formulation of k -NN (Holmes & Adams 2002; Everson & Fieldsend 2004; Manocha & Girolami 2007), which allows for posterior inference over k and the similarity metric.

Finally, one may simply try to find a heuristic, reasonable estimate of the uncertainty of the new data. This is for instance how the photo- z uncertainties in the SDSS database have been computed (Abazajian et al. 2009). Here, a hyperplane was fitted to the nearest 100 neighbours in colour space, and the mean deviations of the redshifts from this hyperplane were found to be good estimates of the errors.

To our knowledge, there is no accepted way of dealing with all uncertainty issues short of Monte Carlo sampling. In this paper, we have therefore chosen to ignore the question relating to uncertainties, focusing solely on demonstrating the performance gain of combining k -NN and feature selection.

2.2 Choosing the number of neighbours

In most versions of k -NN, including the vanilla version, one much choose the number of neighbours, k , to average over. Increasing k implies that a prediction will be based on the average of many samples, which reduces the variance of the classifier but may increase its bias (for a discussion of the bias-variance decomposition of the error of k -NN regression we refer to Hastie et al. 2009). A stan-

dard technique for choosing k is cross-validation (CV). In M -fold CV, the available data S is randomly partitioned into M subsets S_1, \dots, S_M of (almost) equal size. Let $S_{\setminus i} = \bigcup_{j=1, \dots, M, j \neq i} S_j$ denote all data points except those in S_i . For each $i = 1, \dots, M$, an individual model is built by applying the algorithm to the training data $S_{\setminus i}$. This model is then evaluated using the test data in S_i . The average error is called *cross-validation error* and is a predictor of the generalisation performance of the algorithm. To choose k for k -NN using M -fold CV, S is split into M subsets. For each fold $i = 1, \dots, M$ k -NN models are built and tested using different values for k (say, $k = 1, 3, 5, \dots$). The k with the lowest CV error is finally selected.

It must be stressed that the data used for model selection must be independent from data for assessing the final performance of a model.

2.3 Informative features

The use of appropriate features is crucial for machine learning. Standard features in astronomy are, for instance, magnitudes or the derived colours. The performance of a model can, however, often be improved by considering additional features or special combinations of features (thus, effectively changing the underlying distance metric d).¹ We employ automatic *feature selection* to pick the most informative features for our regression task.

2.3.1 Feature selection

The goal of feature selection is to reduce the dimensionality of the input space by selecting the most informative features. A direct way to select such informative features is to systematically try various combinations of features and select the subset with the most promising accuracy for the final model (based on a certain evaluation criterion such as CV). In theory, one would like to try every possible combination of features, but in practice this is often infeasible due to the induced exponential runtime. In the literature, different techniques have been proposed to address this issue such as the idea to maximise the probability of finding the best combination of features. We refer to Guyon & Elisseeff (2003) for an introduction to feature selection.

Standard alternatives to such an exhaustive search are *forward* and *backward feature selection* (Hastie et al. 2009), which aim at selecting informative features in an incremental manner. For the case of forward selection, one starts by selecting the most promising feature by assessing the predictive power of each of the D features. In the second iteration, the first feature is kept and a second one is selected based on the predictive power of both the first *and* the second feature. This process is repeated until the number \bar{d} of desired features is selected. Backward elimination works similarly. However, instead of incrementally adding features, one removes a feature at a time, starting with all D features being selected.

Even forward and backward feature selection are still

¹ For instance, the SDSS pipeline resorts to two different types of magnitudes via the linear model $\text{psfMag} - \text{cModelMag} > 0.145$ to classify photometric objects as ‘galaxy’ or ‘point-like’.




Figure 1. An example spectrum of a galaxy from the SDSS database (black curve) overlaid by the five bandpass filters of SDSS (Fukugita et al. 1996).

computationally demanding, but using clever implementations and data structures one may parallelize the procedure. This paper uses a massively-parallel matrix-based implementation combining incremental feature selection and nearest neighbour models, recently proposed by Gieseke et al. (2014a). For more details on the implementation, we refer to appendix A.

3 EXPERIMENTAL SET-UP

3.1 Data selection

The experiments in this study use photometric data from Sloan Digital Sky Survey (SDSS, York et al. 2000). The data are a subset of SDSS Data Release 7 (DR7, Abazajian et al. 2009), and consist of *psfMag*, *fiberMag*, *petroMag*, *deVMag*, *expMag*, and *modelMag* magnitudes in the *u*, *g*, *r*, *i*, and *z* bands (see Fig. 1) for each galaxy as well as the galaxy’s sSFR and redshift, estimated from spectroscopy. We also include the photometric redshifts estimated by SDSS (Abazajian et al. 2009).

Data are obtained from SDSS CasJobs, using the *SpecPhoto* view, which ensures that objects have clean spectra. Specific star formation rates were taken from Brinchmann et al. (2004)². To clean the data, we apply the following constraints:

- For sSFRs, we require that the estimation was successful (*flag* = 0), and we remove all duplicate galaxies.
- For redshifts, we require that both spectroscopic and photometric estimations were successful (for spectroscopy, *zWarning* = 0; for photometry, *zErr* >= 0).

A sample of 611 479 galaxies meet the above criteria. For a smaller subset of 7799 low-redshift galaxies ($0.0042 < z < 0.33$) within the selected sample, we additionally have photometric sSFR estimations obtained by a template-based modelling approach described in section 3.2. No additional selection criteria have been applied to this subset. In particular, no S/N cut has been used in order to highlight the method’s robustness to varying noise levels. In this work, we do not make use of any S/N information, though special

² We used *specsfrr_avg* from the data located at <http://wwwmpa.mpa-garching.mpg.de/SDSS/DR7/sfrs.html>

treatment of low S/N sources can be incorporated in various ways (see discussion in section 2.1.1).

The experiments will be based on two samples of galaxies: the smaller subset and the full sample, excluding the smaller subset (totalling 603 680 galaxies). The redshift distributions of these two samples can be seen in Fig. 2. The smaller subset consists entirely of low-redshift galaxies, where also the majority of the larger sample can be found. The larger sample consists primarily of galaxies below $z \sim 0.5$, with only a few galaxies at higher redshifts.

3.2 Comparison with other methods

We will compare our results to those of two other methods, one for the sSFR estimations and one for the photo- z estimations.

For the photo- z experiments, we compare our results to the photo- z 's available directly through the SDSS database. These photo- z 's have been estimated using a combination of k -NN (Csabai et al. 2007) and a template-based method (Budavári et al. 2000), as described in Abazajian et al. (2009). The k -NN part of the method differs from our approach in that it bases the estimated photo- z on a local hyperplane fitted to the 100 nearest neighbours, instead of just taking the average (and optimizing the number of neighbours), as we do. It is also important to note that our experimental set-up is different to the one SDSS uses. In particular, SDSS have likely based their photo- z estimates on a much larger training set than we use.

For the sSFR experiments, we compare our estimated sSFRs to those obtained by the standard approach of stellar population synthesis modelling very similar to those used in Gallazzi et al. (2005, 2008); Salim et al. (2007)³. Roughly speaking, a large library of template spectra is generated from stellar population synthesis models. To estimate the SFR of a certain observed galaxy, one would compare the galaxy's spectrum to each of the template spectra. The SFRs of the templates are then weighted based on the likelihood of the template spectra given the real spectrum, resulting in a probability distribution for the SFR. From this distribution, the final SFR of the galaxy is calculated as the expected SFR.

To estimate the sSFR when only photometric information is available, the template spectra are multiplied by the filter transmissions of the particular survey, in our case SDSS (see Fig. 1), to produce template magnitudes. These are then compared to observed ones, and the pipeline described above continues.

3.3 Description of experiments

We considered four different experimental set-ups with the common goal of estimating sSFRs and photo- z 's of galaxies as accurately as possible. The first two experiments were based on the exact same galaxy sample as used for the template-based model (and can thus be compared directly), whereas the last two experiments were based on the total selected galaxy sample mentioned in section 3.1, but with the smaller subset excluded (hereafter referred to as the *larger*

subset). The experiments for sSFR and photo- z estimations were identical in set-up – only the quantity to be estimated changed.

Common to all experiments is that we used the four colours $u - g$, $g - r$, $r - i$, and $i - z$ of the galaxies, and in the experiments with feature selection we additionally included the plain magnitudes u , g , r , i , and z . The magnitudes varied from experiment to experiment, see the summary in Table 1 and detailed description further down. The data and code used for the experiments can be found online, see appendix B.

In each experiment, a nested cross-validation (CV) – an inner and an outer – was used to assess the performance of the k -NN method. Both the inner and outer CV partitioned the data into 10 folds with 9 folds being used for training and the remaining fold being used for testing. For each outer CV, the 9 folds of training data were further partitioned into 10 inner folds for the inner CV. Of these 10 inner folds, 9 were used as training data and the remaining as test data in order to determine the optimal $k \in \{2, 3, 4, \dots, 50\}$, while simultaneously doing feature selection by minimizing the root-mean-square error (RMSE). The exact number of chosen features, as well as which features were chosen, therefore varied across all folds. This simultaneous k determination and feature selection was made possible by the massively parallel GPU implementation of the k -NN algorithm described in Gieseke et al. (2014a). Doing feature selection on the scale of this study is simply not feasible without a highly optimized k -NN implementation.

After the optimal features and optimal k were determined by the inner CV, the performance was assessed by the outer CV.


The performance of each method was therefore assessed 10 times, allowing us to calculate both the means and (population) standard deviations for each of the performance metrics discussed in section 4. As folds in a CV procedure are not fully independent of each other, these standard deviations cannot be interpreted as strict confidence intervals.

To make the estimations by the k -NN, the template-based model (for the sSFR estimations) and the SDSS method (for the photo- z estimations) comparable, the predictions by the latter two methods were divided into the same 10 subsets as used in the outer CV of the k -NN, and the same statistics were calculated. The four experiments were devised as follows:


Experiment 1 The first experiment used the smaller subset (7799 galaxies) and used the four `modelMag` colours $u - g$, $g - r$, $r - i$, and $i - z$ as features. No feature selection was performed, but k was still optimized in each of the inner CV folds. This experiment acts as a baseline for the later feature selection.

Experiment 2 The second experiment again used the smaller subset, but this time all six types of magnitudes (`psfMag`, `fiberMag`, `petroMag`, `devMag`, `expMag`, and `modelMag`) were used. Each type of magnitude gives rise to four colours and five magnitudes, totalling 54 features. A feature selection was performed independently for each outer CV fold to find the best feature combination.

³ J. Brinchmann, private communication.



(a) Small subset.



(b) Entire sample, excluding smaller subset.

Figure 2. Redshift distribution of the two galaxy samples used in the experiments. The entire sample, excluding the smaller subset, additionally contains a single galaxy with $z = -1.93 \times 10^{-4}$, which is not shown in the plot.

Experiment	Sample size	Feature selection	Features
1	7799	No	<code>modelMag</code> ; colours only
2	7799	Yes	<code>psfMag</code> , <code>fiberMag</code> , <code>petroMag</code> , <code>deVMag</code> , <code>expMag</code> , <code>modelMag</code> ; colours and magnitudes
3	603 680	No	<code>modelMag</code> ; colours only
4	603 680	Yes	As selected in experiment 2

Table 1. Summary of experiments. The experiments were based on the four colours $u - g$, $g - r$, $r - i$, and $i - z$, as well as the five magnitudes u , g , r , i , and z , where indicated.

Experiment 3 The third experiment used the larger subset. The features were again only the four `modelMag` colours, and the experiment will serve as a baseline for the k -NN performance on this larger subset.

Experiment 4 The fourth experiment also used the larger subset. The features were chosen to be the overall most informative ones found in experiment 2, based on a median ranking of the importance of each feature across the CV folds. This last experiment will test how well k -NN, with features found from a feature selection on a small data set, can extended to a much larger data set, thus assessing its performance in a ‘big data’ setting.

4 RESULTS AND ANALYSIS

4.1 Specific star formation rate experiments

We evaluate the sSFR experiments using the following performance metrics. In general, we use the logarithm of the ratio of the estimated sSFR to the spectroscopically confirmed, $\Delta\text{sSFR} \equiv \log_{10}(\text{sSFR}_{\text{est}}/\text{sSFR}_{\text{spec}})$. For each CV fold m , we compute the root-mean-square error (RMSE) as

$$\text{RMSE} = \sqrt{\frac{1}{|S_m|} \sum_{n \in S_m} \Delta\text{sSFR}_n^2},$$

where S_m is the test set. We also compute the median of ΔsSFR , as well as the scatter, σ , defined to be the standard deviation of ΔsSFR over the test set. Lastly, we report the fraction of catastrophic outliers, η , defined to be galaxies with $|\Delta\text{sSFR}| > 3\sigma$.

Results of the sSFR experiments can be seen in Table 2. The reported values are the means and standard deviations of each performance metric over the ten CV folds.

Comparing first the results of the experiments on the smaller subset of SDSS (experiment 1 and 2) to the result of the template-based model, we see a clear overall improvement for both experiments. In particular, the median is much improved, showing that k -NN achieves a lower bias.


In addition, doing feature selection (experiment 2) rather than simply using the four `modelMag` colours (experiment 1) further improved the estimations, though not as significant as the differences to the template-based model.

Figure 3 shows the RMSE and standard deviation of the sSFR estimation for each of the ten CV folds of the smaller subset during feature selection (experiment 2). The RMSE and standard deviation are computed each time a feature is added. It is seen that by far the largest gain in accuracy happened with the addition of the first three features (which for all folds are three `modelMag` colours, see below). The error kept decreasing until it was at its lowest at seven to nine added features after which the error started to increase. This is a very commonly seen behaviour for k -NN, and the reason is likely the decreasing quality of the features; as the dimensionality of the feature space increases, we are adding less informative (i.e., noisier) features. The combined effect is that the nearest neighbours to any given data point might change and the estimation will be worse as a result. It is therefore important to stop the feature selection process before the error starts increasing. The results for experiment 2 in Table 2 were achieved in exactly this way, i.e., by stopping when the RMSE was lowest.

To see which features were chosen in experiment 2, and in which order they were chosen, the results for each CV are

Table 2. Root mean square errors (RMSEs), medians and scatter of ΔsSFR , shown as their mean and standard deviations over the ten CV folds for the k -NN regressions and the template-based model.

Experiment	D	RMSE/ $10^{-2} \log(\text{yr}^{-1})$	Median/ $10^{-2} \log(\text{yr}^{-1})$	Scatter, $\sigma/10^{-2} \log(\text{yr}^{-1})$	$\eta/\%$
SDSS subset of 7799 galaxies					
1	4	29.0 ± 1.8	1.63 ± 1.20	28.9 ± 1.7	1.78 ± 0.36
2	8 ^a	27.1 ± 1.5	1.52 ± 0.99	27.0 ± 1.4	1.72 ± 0.32
Template-based model		34.9 ± 1.6	-12.4 ± 0.8	30.4 ± 1.6	3.05 ± 0.35
SDSS subset of 603 680 galaxies					
3	4	30.4 ± 0.6	1.50 ± 0.16	30.4 ± 0.6	1.74 ± 0.05
4	8	28.3 ± 0.4	1.16 ± 0.08	28.3 ± 0.4	1.79 ± 0.04

^a Number of features is the median of the ten CV folds.**Figure 3.** Root mean square error (RMSE) and one standard deviation intervals for each of the ten CV folds of the estimated sSFRs during feature selection. A sharp decrease in error is seen as the first features are added, but it levels off quickly after the first three added features. As features continued to be added the errors started increasing again.

illustrated in Fig. 4. The full list of ranked features can be seen in Fig. C1.

The names of the features are shown to the left of the plot, and the middle ten columns show the ranking of the features for each of the ten CV folds, with a larger bar indicating that the feature was chosen earlier in the feature selection process, and thus has higher importance. A bar is coloured blue if the corresponding feature was selected in the feature selection process. Note that the amount of selected features per fold varies, as do the chosen features themselves. This is due to the differences in the data for each fold. This variation should become less prominent with an increased amount of data, as the folds will statistically become more and more similar. The rightmost column shows the median rank of each feature over all CV folds.

It is seen that the top six features were consistently chosen in each CV fold, except for the u band **psfMag**, which was replaced by the r band in the last fold, see Fig. C1. The remaining chosen features varied more, but were also consistent enough that, except for the r band **psfMag**, no feature below the ninth in the figure was ever chosen. For the overall most informative features (for use in experiment 2 and 4), we chose to select the top eight features from the plot, since eight is the median of the number of chosen features across the CV folds. This is just one particular choice, and

one may equally well use other selection criteria or ranking methods, e.g., ranking based on how often a feature was chosen. Indeed, testing various ranking and selection criteria is an obvious extension to our work, though the exact choices are unlikely to cause significant changes to the results. In summary, we chose to base both ranking and selection on the median.

Returning to the figure, it is interesting that only a single **petroMag** colour was ever selected, even though these are the magnitudes recommended by the SDSS⁴ for use with low-redshift galaxies. Instead, the most prominent features were **modelMag** and **fiberMag** colours, with **modelMag** colours as the top three most informative features. This is not surprising, since the **modelMag** magnitudes are defined as either **expMag** or **deVMag** magnitudes depending on which fits the best.

Interestingly, none of the selected features use the z band, which can likely be explained by the band's low filter transmission, as seen in Fig. 1. This will often result in a low signal-to-noise ratio (S/N). Also interesting is the fact that the u band appears in many of the most informative features, even though it also has a low S/N. The reason is likely that the band captures UV radiation from newly formed stars, thus directly measuring (part) of the SFR.

Looking further down the list of selected features, we see that magnitudes and colours based on **expMag** were generally ranked much higher than their **deVMag** counterparts. This is interesting, as **modelMag**, which dominates the list of informative features, is the better fit of **expMag** and **deVMag**. This could suggest that the **modelMag** mostly resorted to a **deVMag** fit; adding **deVMag** colours (again) would not provide any new information, so the feature selection chooses to add **expMag** colours instead. Indeed, comparing the likelihoods of the **deVMag** and **expMag** fits⁵ reveals that the **deVMag** fit achieved the largest likelihood for $\sim 66\%$ of the galaxies in the smaller subset.

Returning to Table 2 and now considering experiment 3, which used the larger subset, but only the four **modelMag** colours, we see a performance similar to experiment 1, though now with significantly reduced uncertainties due to the larger sample size.

Experiment 4 also used the larger subset, but with the eight colours chosen as the most informative in experiment 2 (the top eight colours in Fig. 4). As noted, the idea behind

⁴ <http://classic.sdss.org/dr7/algorithms/photometry.html>

⁵ Available through the **PhotoObjAll** table in the SDSS database.




Figure 4. Ranking of the top 25 most important features from the feature selection in experiment 2. To the left are the feature names, while the rightmost column shows the median rank of each feature across all CV folds. Each of the other columns shows the feature ranking in a particular CV fold. The larger the bar for a certain feature, the more important the feature was. Blue bars show features that were chosen during the feature selection as the most informative in a particular CV fold. Because of the differences in the data used in each CV fold the exact features selected as important, as well as the number of chosen features per fold, will vary. The number of chosen features vary from 7 through 9, with a median of 8.

this experiment was to see how features selected on a smaller subset generalise to a larger one. This is important to know if this method is to be applied to a larger part of SDSS without any spectroscopically determined sSFRs to check for consistency with. The results from experiment 4 show that the feature selection from experiment 2 did indeed increase the performance of the method compared to using the standard colours (experiment 3). The fact that the results of experiment 4 were consistent with those of experiment 2, shows that the most informative features can indeed be determined from a smaller subset and then used on a larger. Additionally, it shows that k -NN regression can be an effective method for determining sSFRs from photometric data, even when the features are determined from a much smaller subset.

Figure 5 shows the correlations between the spectroscopically determined sSFRs and the corresponding estimations from the template-based model as well as each of the four experiments. Looking at the estimations from the template-based model (Fig. 5a) it is immediately clear where it falls short: it seems to consistently underestimate the sSFRs of the low-sSFR galaxies. The distribution for high-sSFR galaxies also seems slightly skewed towards underestimation.

The estimations done by the k -NN regression (Figs. 5b and 5c) were clearly better than those from the template-based model. The distribution for high-sSFR galaxies seems quite symmetric, while for the low-sSFR galaxies it appears slightly skewed towards overestimating the sSFRs.

The same trends can be seen in the estimations by the k -NN regression on the larger subset (Figs. 5d and 5e); a symmetric mode for the high-sSFR galaxies and a slightly

skewed mode for the low-sSFR galaxies, though not as pronounced as for the smaller subset.

For all k -NN experiments, the distribution at highest sSFRs seem to skew towards underestimation. This is likely due to the inherent inability of k -NN to extrapolate beyond the distribution of the training set; as there are only few data at these sSFRs, it is likely that the average of the nearest neighbours (in colour space) will drive the estimated sSFR towards lower values. Apart from choosing a different method than k -NN, an obvious remedy would be to include more galaxies in the training set to cover more of the colour-magnitude and sSFR space. Another possibility would be to include colours from other surveys, thereby increasing the dimensionality of the colour-magnitude space. This could potentially add the extra information needed in order to move the galaxies closer to others with similar sSFRs. Indeed, [Salim et al. \(2005\)](#) showed that a combination of SDSS and *GALEX* ([Martin et al. 2005](#)) photometry led to a significant improvement in the estimation of SFRs over using just SDSS photometry. It is natural to assume that this would also be the case with our method.

When looking at Fig. 5, all distributions seem to have a hump around $(-11, -12)$, where the sSFRs are somewhat underestimated. It appears as if galaxies from the green valley get mixed up with quenched galaxies. The problem also seems to be present for the template-based model, indicating that there may not be enough information in the SDSS magnitudes to distinguish these galaxies from quenched ones. Giving the galaxies a closer look would be an obvious next step to further increase the accuracy of the methods. It is, however, clear that the k -NN method works equally well for estimating sSFRs for both main-sequence and quenched




Figure 5. Correlations between the estimated and spectroscopically determined sSFRs for the template-based model and the four experiments. The colour coding indicates the amount of galaxies in each bin.

galaxies, which is a rare quality for sSFR estimation methods in general.

Figure 6 shows the sSFR residuals as function of spectroscopic redshift, with the orange line showing the running median of the underlying distribution. The thick bars span the 15.87th through the 84.13th percentile ($\pm 1\sigma$), and the thin bars span the 2.28th through the 97.72th percentile ($\pm 2\sigma$).

The template-based model (Fig. 6a) has a clear tendency to underestimate the sSFR throughout the entire redshift range. Our k -NN model (Figs. 6b and 6c) performs a lot better, with a running median close to 0 at all redshifts. The scatter around the running median seems similar for both models, which is also apparent from Table 2.

Although the data are limited to rather low redshifts, it is reassuring to see that there appears to be no significant increase in either bias or scatter, even at the highest redshifts with our model. Note that the redshift was not part of the features used by our method. Estimation of sSFRs at all redshifts is based solely on colours and magnitudes of the galaxies.

4.2 Redshift experiments

The accuracy of the photo- z experiments is evaluated with the following metrics. We define the normalised redshift estimation error as $\Delta z' = \Delta z / (1 + z)$, where $\Delta z = z_{\text{phot}} - z_{\text{spec}}$. Following Ilbert et al. (2006), we define a catastrophic outlier as a galaxy with $|\Delta z'| > 0.15$ and η as the fraction

of catastrophic outliers in a given experiment. We further use the definition of the normalised median absolute deviation as $\sigma_{\text{NMAD}} = 1.48 \times \text{median}(|\Delta z'|)$. Following Dahlen et al. (2013), we define $\sigma_{\text{RMS}} = \langle \Delta z'^2 \rangle^{1/2}$ and σ_{O} as being the σ_{RMS} after catastrophic outliers have been removed. We also evaluate the bias, given as the mean normalised error, $\text{bias}_z = \langle \Delta z' \rangle$, once again excluding catastrophic outliers.

Table 3 presents the results obtained in the various experiments. The results are calculated by combining the results from the test sets in each of the 10 CV folds.

Considering first the experiments on the smaller subset, the SDSS method is quite consistently outperforming our experiment 1, though the differences are within one standard deviation. Our experiment 2, however, is consistently outperforming the SDSS method, though again the differences are within one standard deviation. Comparing our experiment 1 and 2 shows a much more significant difference; the chosen features clearly outperformed the four standard colours.

Figure 7 shows the 25 most important features obtained from the feature selection in experiment 2, with the features chosen as most informative coloured in blue. The full list of ranked features can be seen in Fig. C2. The top seven features were quite consistently chosen as the most important, whereas the remaining chosen features in each CV fold are a lot more scattered than for sSFR estimation. The number of chosen features also vary much more: from six to eleven features are chosen in the folds. The median number of selected features was 7.5, so the top eight features in Fig. 7 were cho-




Figure 6. sSFR residuals as function of redshift for the two galaxy samples used in the experiments. The colour coding of the distributions indicates the amount of galaxies in each bin. The orange line shows the running median of the underlying distribution, the thick bars span the 15.87th through the 84.13th percentile ($\pm 1\sigma$), and the thin bars span the 2.28th through the 97.72th percentile ($\pm 2\sigma$). Residual plots for experiments 1 and 3 can be found in appendix D.

Table 3. Results from the photo- z estimation experiments. Evaluation metrics are the bias, $\text{bias}_z = \langle \Delta z' \rangle$, the normalised root-mean-square (RMS) error, $\sigma_{\text{RMS}} = \langle \Delta z'^2 \rangle^{1/2}$, the RMS error with outliers removed, σ_O , the normalised median absolute deviation, $\sigma_{\text{NMAD}} = 1.48 \times \text{median}(|\Delta z'|)$, and the fraction of catastrophic outliers, η . The standard deviations shown are calculated over the 10 CV folds.

Experiment	D	$\text{bias}_z / 10^{-4}$	$\sigma_{\text{RMS}} / 10^{-2}$	$\sigma_O / 10^{-2}$	$\sigma_{\text{NMAD}} / 10^{-2}$	$\eta / 10^{-2} \%$
SDSS subset of 7799 galaxies						
1	4	-6.13 ± 9.80	2.19 ± 0.08	2.17 ± 0.07	1.72 ± 0.10	2.56 ± 5.13
2	8 ^a	4.29 ± 7.91	1.83 ± 0.12	1.82 ± 0.10	1.45 ± 0.07	1.28 ± 3.85
SDSS	4 ^b	-5.84 ± 10.60	2.01 ± 0.12	1.99 ± 0.13	1.54 ± 0.09	3.85 ± 5.88
SDSS subset of 603 680 galaxies						
3	4	3.57 ± 0.61	2.29 ± 0.04	2.26 ± 0.04	1.83 ± 0.03	3.56 ± 0.79
4	8	3.00 ± 0.56	1.74 ± 0.03	1.73 ± 0.03	1.40 ± 0.02	0.795 ± 0.232
SDSS	4 ^b	5.29 ± 5.38	2.22 ± 0.02	2.12 ± 0.02	1.65 ± 0.02	8.58 ± 0.89

^a Number of features is the median of the ten CV folds.

^b SDSS additionally fitted a hyperplane in order to make estimations.

sen as basis for experiment 4. The varying features as well as the number of chosen features for each CV fold can be an indication that many magnitudes and colours have very similar information content. Thus, even small differences in the datasets used in each CV fold can be enough to change the features deemed most informative. Using a larger dataset for the feature selection will likely make the chosen features more stable.

It is interesting to see that, while three of the four **modelMag** colours are among the selected features, they are not the most informative. The **fiberMag** colours appear to contain more information for photo- z estimation.

Another interesting observation is that the z band **expMag** was chosen consistently in all but one CV folds. In the fourth fold, the z band **deVMag** was chosen instead of the **expMag**. Having a single measure of the z band magnitude therefore seems to be important for photo- z estimation. This is rather surprising, given the z band's low S/N and the fact that all galaxies in the small subset have $z \lesssim 0.33$.

Returning again to Table 3, it is expected that the SDSS method outperform our experiment 1. Even though we use the same features, the SDSS estimate uses a hyperplane fit to the nearest 100 samples. This will act as regularisation, making estimations less susceptible to outliers.

Considering now the experiments on the larger subset, the results are qualitatively as before, but with significantly

reduced error bars. Overall, SDSS outperforms our experiment 3, which again uses the same features. As before, this is to be expected. Interestingly, our experiment 3 has a significantly lower outlier rate η than SDSS, but that is likely due to SDSS training on a larger sample. As k -NN is unable to extrapolate beyond just the mean of the nearest data points, our method will not be able to estimate a redshift outside the redshift range of the training set. As SDSS has likely used a much larger training set with more high- z galaxies, it is plausible that their method has been ‘confused’ by galaxies with similar colours, but much higher redshifts, leading to a significantly larger redshift estimations, and thus the possibility of more outliers.

Experiment 4 significantly outperformed both our experiment 3 and, more interestingly, the photo- z estimations from SDSS. We are thus able to achieve much better performance by using optimal features instead of the standard ones, even when using additional modelling as SDSS does.

Figure 8 shows correlations between estimated photo- z and the spectroscopically derived redshift. The spectroscopically determined redshifts have a sharp cut around $z \sim 0.33$, after which there are only few galaxies. This is a result of our data selection.

Figures 8a and 8d show the photo- z estimations done by SDSS, i.e., not by our model. Figures 8b and 8e show the photo- z estimations done using our k -NN method, but




Figure 7. Ranking of the 25 most important features according the feature selection in experiment 2. To the left are the feature names, while the rightmost column shows the median rank of each feature across all CV folds. Each of the other columns shows the feature ranking in a particular CV fold. The larger the bar for a certain feature, the more important the feature was. Blue bars show features that were chosen during the feature selection as the most informative in a particular CV fold. Because of the differences in the data used in each CV fold the exact features selected as important, as well as the number of chosen features per fold, will vary. The number of chosen features vary from 6 through 11 with a median of 7.5.




Figure 8. Correlations between the estimated photo- z and spectroscopically determined z for the SDSS photo- z method and our k -NN method. The colour coding of the distributions indicates the amount of galaxies in each bin.

using only the four `modelMag` colours. Finally, Figs. 8c and 8f show the photo- z estimations done using our k -NN method, including the feature selection.

Focusing first on the experiments using the small subset, we see that distributions resulting from our k -NN method and SDSS's are, qualitatively, quite similar. Experiment 2, which used feature selection, seems to have a slightly more symmetric distribution around the diagonal and appears to work better at the smallest redshifts, but is otherwise very similar to the other two experiments. Again, it is worth noting that the SDSS estimations have likely used a much larger training set than just this subset, which our model is restricted to. The fact that the results are so similar shows that the k -NN method does not need a large training sample to produce accurate estimations.

Turning now to the experiments using the large subset, the SDSS method appears to result in more extreme outliers than ours. This observation may, however, be misleading. As the SDSS estimations are likely based on a much larger training set including galaxies with much higher and lower redshifts, estimations outside the distribution of our subset are entirely possible. Our training set is limited to low redshifts, meaning we cannot estimate redshifts outside this range. Thus, our estimations have a clear cut at a redshift similar to that of the spectroscopic redshift cut, with only a few outliers due to a few high redshift galaxies.

Ignoring for a moment estimations above $z_{\text{phot}} \sim 0.33$, the SDSS estimation seems to perform better than our experiment 3 (Fig. 8e), which only used the four `modelMag` colours. The SDSS photo- z distribution seems tighter around the diagonal, which is likely a result of the hyperplane fit acting as regularisation. Both methods do, however, significantly overestimate at the lowest redshifts.

Figure 8f shows the estimations done by our k -NN method, using the features obtained from the feature selection process in experiment 2. Compared to Fig. 8e, there is less scatter and the distribution is significantly tighter around the diagonal. Comparing with the SDSS estimations (Fig. 8d), we perform significantly better at the lowest redshifts, with the added bonus of a more symmetric distribution.

Finally, Fig. 9 shows the photo- z residuals as function of spectroscopic redshift. The rather sharp slopes at $z \sim 0.3$ for our estimations are a result of the cut in the spectroscopic redshift as discussed previously. The SDSS estimations do not exhibit this slope due to the larger training set, making extrapolation beyond $z \sim 0.3$ much more likely.

Considering the photo- z experiments on the small subset, there is not much difference between the estimates from SDSS (Fig. 9a) and those from our k -NN method (Fig. 9c). Note that the estimations from our method have been obtained using feature selection. Both estimation methods appear to overestimate the redshift at low redshifts, though it is more pronounced for the SDSS method. At higher redshifts, the methods both slightly underestimate the redshifts. At the highest redshifts the SDSS method appears to overestimate slightly, while our k -NN method seems to underestimate the redshift. This underestimation is a consequence of the slope, as the training set used for our method contains only a few galaxies with $z \gtrsim 0.33$. Therefore, one should not conclude too much from this underestimation.

The picture is very similar when considering the exper-

iments on the large subset (Figs. 9b and 9d). There is a tendency to overestimate the redshift at small z and underestimate it at higher z . For both experiments, however, the median residual is always close to zero. There are a few extra galaxies at $z > 0.5$ not shown in these plots, in order to keep the main galaxy sample detailed. Both methods significantly underestimate the redshifts of these high- z galaxies with roughly the same amount.

From the plots in Fig. 9, it is clear that using just the most important features, we can achieve a similar performance to fitting a hyperplane to the nearest neighbours, though at a much lower computational cost once the features have been determined.

5 DISCUSSION AND CONCLUSIONS

In the coming years, increasingly larger astronomical surveys will produce unprecedented amounts of data. Many of these data will require accurate estimations in near real-time, which is not feasible with traditional methods. Machine learning is well-suited to address this challenge.

This work has exemplified this by showing how machine learning can be used to not only estimate specific star formation rates (sSFRs) and photometric redshifts (photo- z 's) of galaxies, but also to identify the most informative features for these tasks, thereby increasing accuracy further. We have shown how the simple, yet powerful non-parametric k nearest neighbours (k -NN) method significantly outperforms the traditional method of simulated template spectra for estimating sSFRs, achieving a RMSE of $(2.90 \pm 0.18) \times 10^{-1} \log(\text{yr}^{-1})$ (the \pm values refer to the standard deviation over the non-independent CV folds) compared to a template-based method's $(3.49 \pm 0.16) \times 10^{-1} \log(\text{yr}^{-1})$, when using the exact same input features. Adding a *feature selection* to the k -NN method increased its performance, achieving a RMSE of $(2.71 \pm 0.15) \times 10^{-1} \log(\text{yr}^{-1})$. Similarly, the fraction of catastrophic outliers reduced from the template-based method's $(3.05 \pm 0.35)\%$ to $(1.72 \pm 0.32)\%$, when using k -NN and feature selection.

We see a similar pattern when considering photo- z estimation. Here, the k -NN method achieves a normalised median absolute deviation of $(1.72 \pm 0.10) \times 10^{-2}$, which reduces to $(1.45 \pm 0.07) \times 10^{-2}$ when doing feature selection, compared to $(1.54 \pm 0.09) \times 10^{-2}$ achieved by SDSS. The method used by SDSS even included a hyperplane fit and while that improves estimations, it also significantly increases the required amount of computations per estimate.

Applying the k -NN method to a larger subset of SDSS of 603 680 galaxies, we achieve a RMSE of $(3.04 \pm 0.06) \times 10^{-1} \log(\text{yr}^{-1})$ for sSFR estimation, when using the same four features as the template-based method. By using the features selected in the feature selection on the smaller subset, we are able to decrease the error further to $(2.83 \pm 0.04) \times 10^{-1} \log(\text{yr}^{-1})$. For photo- z estimation, we achieve a normalised median absolute deviation of $(1.83 \pm 0.03) \times 10^{-2}$, which reduces to $(1.40 \pm 0.02) \times 10^{-2}$ when doing feature selection, compared to $(1.65 \pm 0.02) \times 10^{-2}$ achieved by SDSS. This shows that not only can features selected for a smaller subset be directly transferred to a much larger one yielding similar per-




Figure 9. Redshift residuals as function of redshift for the two galaxy samples used in the experiments. The colour coding of the distributions indicates the amount of galaxies in each bin. The orange line shows the running median of the underlying distribution, the thick bars span the 15.87th through the 84.13th percentile ($\pm 1\sigma$), and the thin bars span the 2.28th through the 97.72th percentile ($\pm 2\sigma$). The sharp slopes seen in (c) and (d) are a consequence of the training set containing only few galaxies with $z \gtrsim 0.33$. As the k -NN method is not well suited for extrapolation, only few galaxies will have an estimated photo- $z \gtrsim 0.33$. Residual plots for experiments 1 and 3 can be found in appendix D.

formance, the estimations done by the selected features can even significantly outperform more computationally intensive modelling.

An advantage of a template-based method is the gain in physical knowledge from the simulations. The feature selection for the k -NN method can provide hints to which features contain the most information, but a deeper understanding of why these particular features contain more information requires further investigation and is outside the scope of this work. The k -NN method does, however, have advantages over a template-based method in that it is faster and will not be prone to errors resulting from approximations or wrong assumptions done in the model building process. This study shows that machine learning methods, here exemplified by k -NN regression, should be considered a viable alternative to the traditional template-based method in situations where high accuracy or computational efficiency is required. In particular, adding a feature selection step to the machine learning methods, instead of relying on traditionally used features, should be considered part of the standard toolbox.

ACKNOWLEDGEMENTS

We sincerely thank Jarle Brinchmann for providing us with photometric estimations of masses and star formation rates (through private communication), and for the spectroscopic SFRs made available at <http://wwwmpa.mpa-garching.mpg.de/SDSS/DR7/sfrs.html>. We also thank the SDSS collaboration for making their reduced data available.

This research made use of NASA’s Astrophysics Data System; NumPy & SciPy (Jones et al. 01 ; van der Walt et al. 2011); the IPython package (Perez & Granger 2007); Scikit-learn (Pedregosa et al. 2011); Pandas (McKinney 2010); matplotlib, a Python library for publication quality graphics (Hunter 2007); Seaborn (Waskom et al. 2016).

KSS, CI, and KSP gratefully acknowledge support from The Danish Council for Independent Research | Natural Sciences through the project “Surveying the sky using machine learning”.

Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Soci-

ety, and the Higher Education Funding Council for England. The SDSS Web Site is <http://www.sdss.org/>.

The SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions. The Participating Institutions are the American Museum of Natural History, Astrophysical Institute Potsdam, University of Basel, University of Cambridge, Case Western Reserve University, University of Chicago, Drexel University, Fermilab, the Institute for Advanced Study, the Japan Participation Group, Johns Hopkins University, the Joint Institute for Nuclear Astrophysics, the Kavli Institute for Particle Astrophysics and Cosmology, the Korean Scientist Group, the Chinese Academy of Sciences (LAMOST), Los Alamos National Laboratory, the Max-Planck-Institute for Astronomy (MPIA), the Max-Planck-Institute for Astrophysics (MPA), New Mexico State University, Ohio State University, University of Pittsburgh, University of Portsmouth, Princeton University, the United States Naval Observatory, and the University of Washington.

REFERENCES

- Abazajian K. N., et al., 2009, *ApJS*, **182**, 543
- Abdalla F. B., Banerji M., Lahav O., Rashkov V., 2011, *MNRAS*, **417**, 1891
- Almosallam I. A., Lindsay S. N., Jarvis M. J., Roberts S. J., 2016, *MNRAS*, **455**, 2387
- Altman N. S., 1992, *Am. Stat.*, **46**, 175
- Arnouts S., et al., 2013, *A&A*, **558**, A67
- Balogh M. L., et al., 2016, *MNRAS*, **456**, 4364
- Benítez N., 2000, *ApJ*, **536**, 571
- Bentley J. L., 1975, *Commun. ACM*, **18**, 509
- Bolzonella M., Miralles J.-M., Pelló R., 2000, *A&A*, **363**, 476
- Brammer G. B., van Dokkum P. G., Coppi P., 2008, *ApJ*, **686**, 1503
- Brinchmann J., Charlot S., White S. D. M., Tremonti C., Kauffmann G., Heckman T., Brinkmann J., 2004, *MNRAS*, **351**, 1151
- Budavári T., Szalay A. S., Connolly A. J., Csabai I., Dickinson M., 2000, *AJ*, **120**, 1588
- Carnero A., Sánchez E., Croce M., Cabré A., Gaztañaga E., 2012, *MNRAS*, **419**, 1689
- Carrasco Kind M., Brunner R. J., 2013, *MNRAS*, **432**, 1483
- Cayton L., 2012, in Proc. IEEE Int. Parallel Distrib. Processing Symp., IEEE, pp 402–413, [doi:10.1109/ipdps.2012.45](https://doi.org/10.1109/ipdps.2012.45)
- Charlot S., Kauffmann G., Longhetti M., Tresse L., White S. D. M., Maddox S. J., Fall S. M., 2002, *MNRAS*, **330**, 876
- Chen Y.-M., Wild V., Kauffmann G., Blaizot J., Davis M., Noeske K., Wang J.-M., Willmer C., 2009, *MNRAS*, **393**, 406
- Collister A. A., Lahav O., 2004, *PASP*, **116**, 345
- Conroy C., 2013, *ARA&A*, **51**, 393
- Csabai I., Dobos L., Trencsényi M., Herczegh G., Józsa P., Purger N., Budavári T., Szalay A. S., 2007, *Astron. Nachr.*, **328**, 852
- D’Isanto A., Cavauoti S., Brescia M., Donalek C., Longo G., Riccio G., Djorgovski S. G., 2016, *MNRAS*, **457**, 3119
- Dahlen T., et al., 2013, *ApJ*, **775**, 93
- Everson R. M., Fieldsend J. E., 2004, in , Lecture Notes in Computer Science. Springer Science + Business Media, pp 654–659, [doi:10.1007/978-3-540-28651-6_96](https://doi.org/10.1007/978-3-540-28651-6_96), http://dx.doi.org/10.1007/978-3-540-28651-6_96
- Fillingham S. P., Cooper M. C., Wheeler C., Garrison-Kimmel S., Boylan-Kolchin M., Bullock J. S., 2015, *MNRAS*, **454**, 2039
- Fukugita M., Ichikawa T., Gunn J. E., Doi M., Shimasaku K., Schneider D. P., 1996, *AJ*, **111**, 1748
- Gallazzi A., Charlot S., Brinchmann J., White S. D. M., Tremonti C. A., 2005, *MNRAS*, **362**, 41
- Gallazzi A., Brinchmann J., Charlot S., White S. D. M., 2008, *MNRAS*, **383**, 1439
- Garcia V., Debreuve E., Nielsen F., Barlaud M., 2010, in Proc. IEEE Int. Conf. Image Process., IEEE, pp 3757–3760, [doi:10.1109/icip.2010.5654017](https://doi.org/10.1109/icip.2010.5654017)
- Geach J. E., 2012, *MNRAS*, **419**, 2633
- Gieseke F., Posterer K. L., Oancea C., Igel C., 2014a, in Proc. Eur. Symp. Artificial Neural Networks, Comput. Intell. Mach. Learn. (ESANN). pp 87–92
- Gieseke F., Heinemann J., Oancea C., Igel C., 2014b, in Proc. Int. Conf. Mach. Learn. (ICML). pp 172–180
- Graham M. J., Djorgovski S. G., Mahabal A. A., Donalek C., Drake A. J., 2013, *MNRAS*, **431**, 2371
- Guyon I., Elisseeff A., 2003, *J. Mach. Learn. Res.*, **3**, 1157
- Hastie T., Tibshirani R., Friedman J., 2009, The Elements of Statistical Learning, second edn. Springer Series in Statistics, Springer New York, New York, NY, [doi:10.1007/978-0-387-84858-7](https://doi.org/10.1007/978-0-387-84858-7)
- Hildebrandt H., et al., 2010, *A&A*, **523**, A31
- Ho S., et al., 2012, *ApJ*, **761**, 14
- Holmes C. C., Adams N. M., 2002, *J. R. Stat. Soc. Ser. B Stat. Methodol.*, **64**, 295
- Hoyle B., Rau M. M., Zitlau R., Seitz S., Weller J., 2015, *MNRAS*, **449**, 1275
- Hunter J. D., 2007, *Comput. Sci. Eng.*, **9**, 90
- Ilbert O., et al., 2006, *A&A*, **457**, 841
- Indyk P., Motwani R., 1998, in Proc. 30th Annu. ACM Symp. Theory of Computing. STOC ’98. ACM, New York, NY, USA, pp 604–613, [doi:10.1145/276698.276876](https://doi.org/10.1145/276698.276876)
- James G., Witten D., Hastie T., Tibshirani R., 2013, An Introduction to Statistical Learning. Springer Texts in Statistics Vol. 103, Springer New York, New York, NY, [doi:10.1007/978-1-4614-7138-7](https://doi.org/10.1007/978-1-4614-7138-7)
- Jones E., Oliphant T., Peterson P., et al., 2001–, SciPy: Open source scientific tools for Python
- Kauffmann G., et al., 2003, *MNRAS*, **341**, 33
- Kennicutt Jr. R. C., 1998, *ARA&A*, **36**, 189
- Kennicutt R. C., Evans N. J., 2012, *ARA&A*, **50**, 531
- Kremer J., Gieseke F., Pedersen K. S., Igel C., 2015, *Astron. Comput.*, **12**, 67
- Kügler S. D., Polsterer K., Hoecker M., 2015, *A&A*, **576**, A132
- Li L., Zhang Y., Zhao Y., 2008, *Sci. China Ser. G Phys. Mech. Aston.*, **51**, 916
- Manocha S., Girolami M., 2007, *Pattern Recognit. Lett.*, **28**, 1818
- Maraston C., Pforr J., Renzini A., Daddi E., Dickinson M., Cimatti A., Tonini C., 2010, *MNRAS*, **407**, 830
- Martin D. C., et al., 2005, *ApJ*, **619**, L1
- McKinney W., 2010, in van der Walt S., Millman J., eds, Proc. 9th Python in Sci. Conf. pp 51 – 56
- Nakasato N., 2012, *J. Comput. Sci.*, **3**, 132
- Pacifici C., et al., 2015, *MNRAS*, **447**, 786
- Padmanabhan N., et al., 2007, *MNRAS*, **378**, 852
- Pedregosa F., et al., 2011, *J. Mach. Learn. Res.*, **12**, 2825
- Perez F., Granger B. E., 2007, *Comput. Sci. Eng.*, **9**, 21
- Polsterer K. L., Zinn P.-C., Gieseke F., 2013, *MNRAS*, **428**, 226
- Polsterer K. L., Gieseke F., Igel C., Goto T., 2014, in Manset N., Forshay P., eds, ASP Conf. Ser. Vol. 485, Astron. Data Analysis Software and Systems XXIII. p. 425
- Richards J. W., Freeman P. E., Lee A. B., Schafer C. M., 2009, *MNRAS*, **399**, 1044
- Rimoldini L., et al., 2012, *MNRAS*, **427**, 2917
- Salim S., et al., 2005, *ApJ*, **619**, L39
- Salim S., et al., 2007, *ApJS*, **173**, 267
- Sánchez C., et al., 2014, *MNRAS*, **445**, 1482
- Smith D. J. B., Hayward C. C., 2015, *MNRAS*, **453**, 1597

- Solorio T., Fuentes O., Terlevich R., Terlevich E., 2005, *MNRAS*, **363**, 543
- Stensbo-Smidt K., Igel C., Zirm A., Pedersen K. S., 2013, in *Proc. IEEE Int. Conf. Big Data*. IEEE, pp 141–144, doi:10.1109/bigdata.2013.6691746
- Tsalmantza P., et al., 2007, *A&A*, **470**, 761
- Walcher J., Groves B., Budavári T., Dale D., 2011, *Ap&SS*, **331**, 1
- Waskom M., et al., 2016, seaborn: v0.7.0 (January 2016), doi:10.5281/zenodo.45133
- Weinberger K. Q., Saul L. K., 2009, *J. Mach. Learn. Res.*, **10**, 207
- Wetzel A. R., Tinker J. L., Conroy C., 2012, *MNRAS*, **424**, 232
- Wetzel A. R., Tinker J. L., Conroy C., van den Bosch F. C., 2013, *MNRAS*, **432**, 336
- Wetzel A. R., Tollerud E. J., Weisz D. R., 2015, *ApJ*, **808**, L27
- Wheeler C., Phillips J. I., Cooper M. C., Boylan-Kolchin M., Bullock J. S., 2014, *MNRAS*, **442**, 1396
- Williams R. J., Quadri R. F., Franx M., van Dokkum P., Labbé I., 2009, *ApJ*, **691**, 1879
- Wuyts S., et al., 2011, *ApJ*, **738**, 106
- Wuyts S., et al., 2013, *ApJ*, **779**, 135
- Xu X., Ho S., Trac H., Schneider J., Poczós B., Ntampaka M., 2013, *ApJ*, **772**, 147
- York D. G., et al., 2000, *AJ*, **120**, 1579
- van der Walt S., Colbert S. C., Varoquaux G., 2011, *Comput. Sci. Eng.*, **13**, 22

APPENDIX A: MASSIVELY PARALLEL GREEDY FEATURE SELECTION

While greedy procedures such as forward or backward feature selection are significantly faster than the exhaustive search for the best-performing features, they can still be very time-consuming, even on training sets of moderate sizes. One way to accelerate such a feature selection step is to speed up the involved nearest neighbour computations. In the literature, various techniques can be found for this task. Typical methods are *k-d trees* (Bentley 1975) or *locality-sensitive hashing* (Indyk & Motwani 1998). However, such tools either perform poorly in higher dimensions or only yield approximate answers. A recent trend in data analytics is to resort to (exact) parallel implementations for many-core devices such as today’s *graphics processing units* (GPUs). For instance, Garcia et al. (2010) make use of highly-tuned GPU matrix multiplication libraries for nearest neighbour search. Other schemes are based on, e.g., adapted spatial search structures (Cayton 2012; Gieseke et al. 2014b; Nakasato 2012).

For the work at hand, we make use of a massively-parallel matrix-based implementation that addresses incremental feature selection and nearest neighbour models recently proposed by Gieseke et al. (2014a). For the sake of completeness, we briefly outline the general workflow of the implementation: The general workflow for the case of forward selection is sketched in Algorithm 1. For a given training set S of labelled samples, start with an empty distance matrix $\mathbf{M} \in \mathbb{R}^{N \times N}$ that contains the current distances between all pairs of training samples. Further, the array `selected_dimensions` indicating the selected features and the array `val_errors` are initialized. The forward feature selection process starts in Step 4: The procedure `GETVALIDATIONERRORS` computes, for each dimension j that has not yet been selected (i.e., `selected_dimensions[j]=0`), the cross-validation error for the case of dimension j being

Algorithm 1 FORWARDSELECTION(S, \bar{d})

Require: Training set $S = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N)\} \subset \mathbb{R}^D \times \mathbb{R}$ and a number $\bar{d} < D$ of desired features.

Ensure: Array `selected_dimensions` with selected features.

- 1: Initialize empty distance matrix $\mathbf{M} \in \mathbb{R}^{N \times N}$;
- 2: `int selected_dimensions[D] = {0, ..., 0}`;
- 3: `float val_errors[D]`;
- 4: **for** $i = 1, \dots, \bar{d}$ **do**
- 5: `val_errors = GETVALIDATIONERRORS(M)`;
- 6: $i_{\min} = \text{GETMINDIM}(\text{val_errors})$;
- 7: `selected_dimensions[i_min] = 1`;
- 8: $\mathbf{M} = \mathbf{M} + \mathbf{M}^{i_{\min}}$;
- 9: **end for**
- 10: **return** `selected_dimensions`

‘added’ to the current set of features. These values are stored in the array `val_errors` and the procedure `GETMINDIM` returns the index of the smallest error contained in it (thus, i_{\min} corresponds to the dimension whose addition leads to the smallest cross-validation error). Afterwards, both `selected_dimensions` and \mathbf{M} are updated accordingly, where $\mathbf{M}^{i_{\min}}$ denotes the all-pairs distance matrix based on dimension i_{\min} only.

The procedure `GETVALIDATIONERRORS` returns the validation errors for all dimensions that have not yet been selected and contributes most to the overall runtime. For each such dimension j , it computes a matrix $\widehat{\mathbf{M}} = \mathbf{M} + \mathbf{M}^j$ containing all pairwise distances with the distances of dimension j being ‘added on the fly’ to the distances that correspond to the previously selected dimensions. This intermediate training set is then used to compute the cross-validation error for the currently selected set of dimensions. It turns out that this procedure and the overall workflow is particularly well-suited for a massively-parallel implementation. Basically, one can parallelise the search over all dimensions that have not yet been selected as well as over the computations of the induced cross-validation errors. By using a standard GPU device, one can reduce the runtime by a factor of up to 150 compared to single-core CPU implementation, hence, reducing the practical runtime needed from hours to minutes only. We refer to Gieseke et al. (2014a) for the technical details and an experimental analysis of the runtimes for typical astronomical data sets.

APPENDIX B: OBTAINING CODE AND DATA

We want to make the results presented in this paper as reproducible as possible, so we are releasing the code and the data obtained from SDSS. The code for the GPU implementation of the nearest neighbours search is available at GitHub: <https://github.com/gieseke/speedynn>. The scripts and data for reproducing the main results of this paper can be found at <http://image.diku.dk/kstensbo/papers/1606.01/>. The page contains a step-by-step guide to setting up the software and recreating the main results presented in this paper.

APPENDIX C: RESULTS FROM FEATURE SELECTION

Figure C1 shows the full feature ranking for the sSFR estimation done in experiment 2.

Figure C2 shows the full feature ranking for the photo- z estimation done in experiment 2.

APPENDIX D: RESIDUAL PLOTS

Residual plots for sSFR experiment 1 and 3 can be seen in Fig. D1 together with residuals of the template-based model, for comparison.

Residual plots for photo- z 's experiment 1 and 3 can be seen in Fig. D2 together with residuals of the SDSS method for the same datasets, for comparison.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.

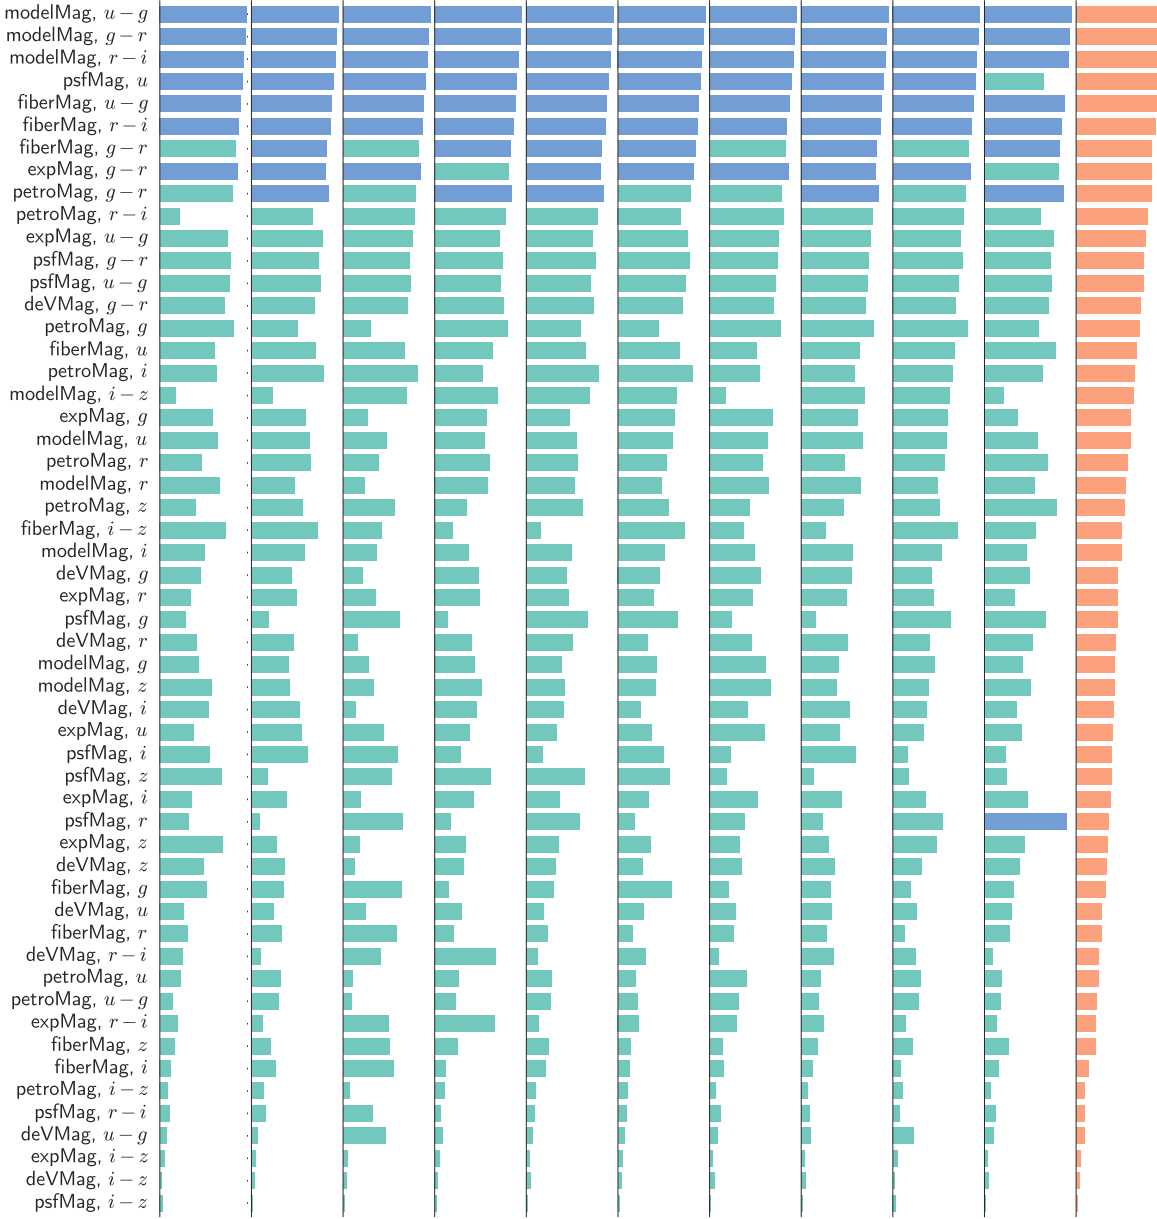


Figure C1. Ranking of features for sSFR estimation according the feature selection in experiment 2. To the left are the feature names, while the rightmost column shows the median rank of each feature across all CV folds. Each of the other columns shows the feature ranking in a particular CV fold. The larger the bar for a certain feature, the more important the feature was. Blue bars show features that were picked out during the feature selection as the most informative in a particular CV fold. Because of the differences in the data used in each CV fold the exact features picked out as important, as well as the number of chosen features per fold, will vary. The number of chosen features vary between 7 and 9 with a median of 8.

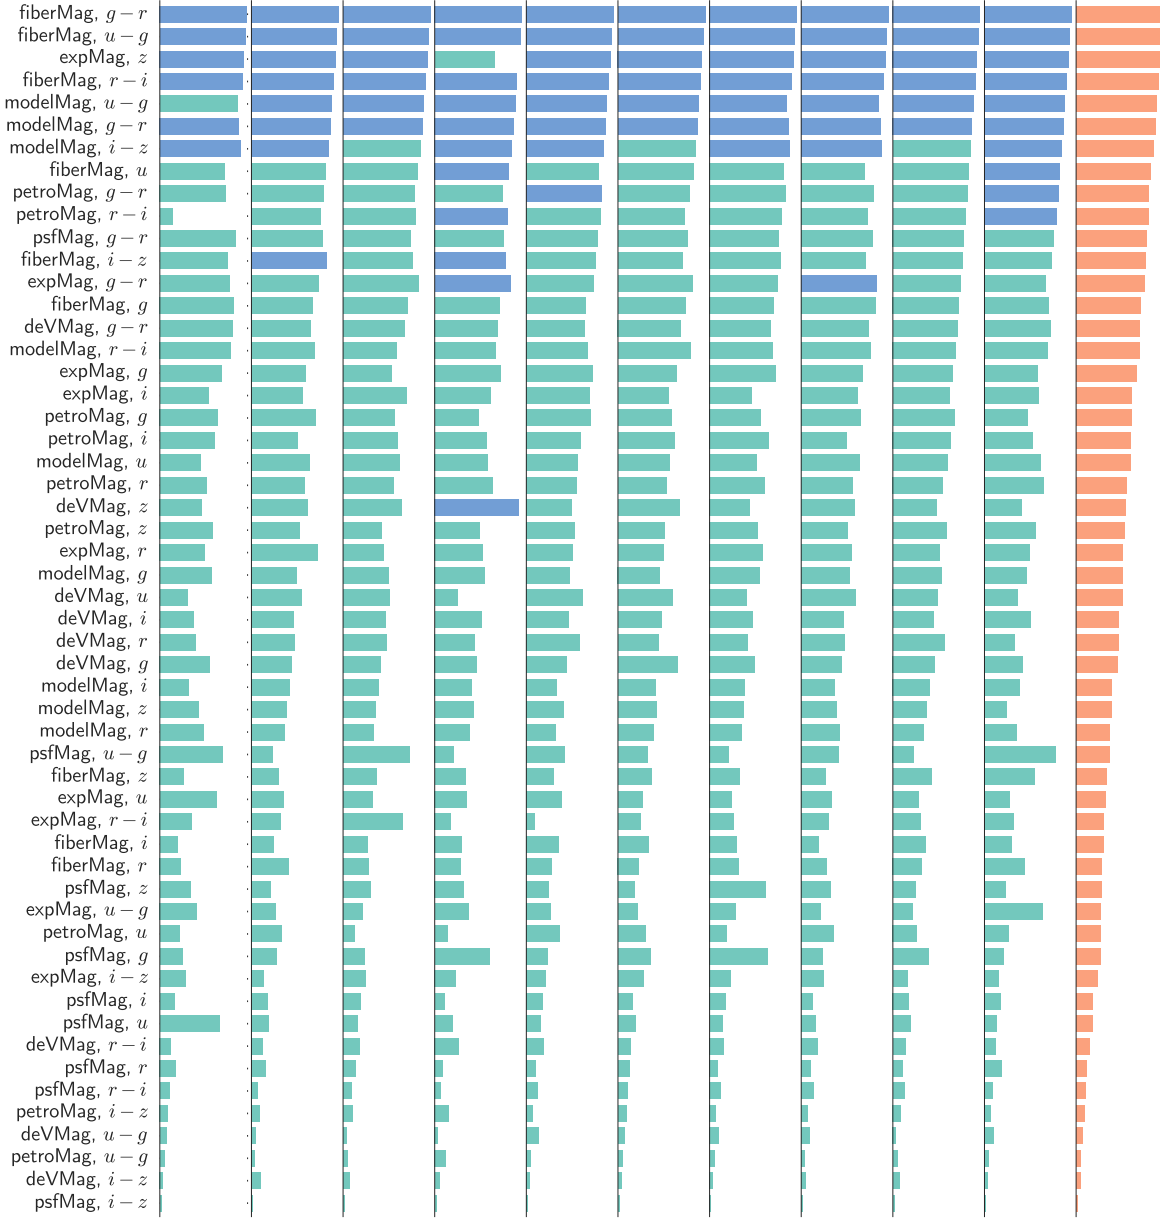


Figure C2. Ranking of features for photo-z estimation according the feature selection in experiment 2. To the left are the feature names, while the rightmost column shows the median rank of each feature across all CV folds. Each of the other columns shows the feature ranking in a particular CV fold. The larger the bar for a certain feature, the more important the feature was. Blue bars show features that were picked out during the feature selection as the most informative in a particular CV fold. Because of the differences in the data used in each CV fold the exact features picked out as important, as well as the number of chosen features per fold, will vary. The number of chosen features vary between 6 and 11 with a median of 8.




Figure D1. sSFR residuals as function of redshift for the two galaxy samples used in the experiments. The colour coding of the distributions indicates the amount of galaxies in each bin. The orange line shows the running median of the underlying distribution, the thick bars span the 15.87th through the 84.13th percentile ($\pm 1\sigma$), and the thin bars span the 2.28th through the 97.72th percentile ($\pm 2\sigma$).

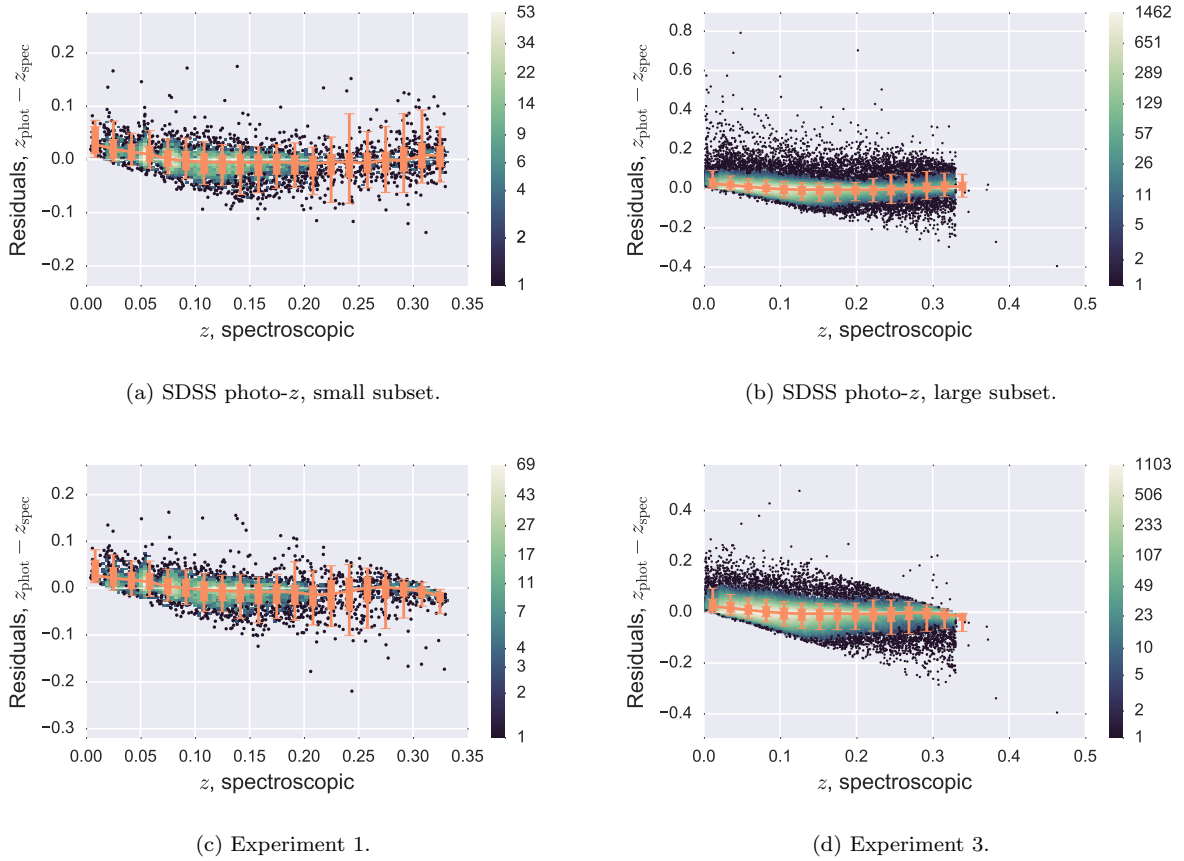


Figure D2. Redshift residuals as function of redshift for the two galaxy samples used in the experiments. The colour coding of the distributions indicates the amount of galaxies in each bin. The orange line shows the running median of the underlying distribution, the thick bars span the 15.87th through the 84.13th percentile ($\pm 1\sigma$), and the thin bars span the 2.28th through the 97.72th percentile ($\pm 2\sigma$). The sharp slopes seen in (c) and (d) are a consequence of the training set containing only few galaxies with $z \gtrsim 0.33$. As the k -NN method is not well suited for extrapolation, only few galaxies will have an estimated photo- $z \gtrsim 0.33$.